

Distributed Control of Generation in a Transmission Grid with a High Penetration of Renewables

Krishnamurthy Dvijotham Dept. of Computer Science and Engineering
University of Washington
Seattle, WA, 98195

Email: dvij at cs.washington.edu

Scott Backhaus MPA Division Los Alamos National Laboratory
Los Alamos, NM 87545

Email: backhaus at lanl.gov Michael Chertkov Theory Division & CNLS
Los Alamos National Laboratory
Los Alamos, NM 87545
Email: chertkov at lanl.gov

November 20, 2012

Abstract

Deviations of grid frequency from the nominal frequency are an indicator of the global imbalance between generation and load. Two types of control, a distributed proportional control and a centralized integral control, are currently used to keep frequency deviations small. Although generation-load imbalance can be very localized, both controls primarily rely on frequency deviation as their input. The time scales of control require the outputs of the centralized integral control to be communicated to distant generators every few seconds. We reconsider this control/communication architecture and suggest a hybrid approach that utilizes parameterized feedback policies that can be implemented in a fully distributed manner because the inputs to these policies are local observables at each generator. Using an ensemble of forecasts of load and time-intermittent generation representative of possible future scenarios, we perform a centralized off-line stochastic optimization to select the generator-specific feedback parameters. These parameters need only be communicated to generators once per control period (60 minutes in our simulations). We show that inclusion of local power

flows as feedback inputs is crucial and reduces frequency deviations by a factor of ten. We demonstrate our control on a detailed transmission model of the Bonneville Power Administration (BPA). Our findings suggest that a smart automatic and distributed control, relying on advanced off-line and system-wide computations communicated to controlled generators infrequently, may be a viable control and communication architecture solution. This architecture is suitable for a future situation when generation-load imbalances are expected to grow because of increased penetration of time-intermittent generation.

1 Problem Setup and Brief Statement of Results

In today's power systems, the system operator performs an Optimal Power Flow (OPF) dispatch periodically with typical time interval being 5, 15, or 60 minutes depending on the Balancing Area [1]. The OPF sets the power outputs of the committed generation to match power demand and minimize generation cost while respecting the capacity limits on lines, ramping constraints and limits

on generators and sometimes taking into account the N-1 security constraints. In between two successive OPFs, the system is automatically controlled by a combination of two mechanisms. The faster of the two, acting on the scale of seconds, is primary frequency control—a fully distributed proportional feedback on locally-measured frequency deviations that may also include a deadband. The slower mechanism, acting on the scale of minutes, is automatic generation control (AGC), also called secondary control—a centralized feedback on the integral of a weighted sum of a centrally measured frequency and tie line flows to neighboring balancing areas.[2]

These combined controls correct deviations in the generation-load balance driven by fluctuations in loads, renewables and other disturbances in the system. However, these mechanisms do not explicitly incorporate line-flow limits, generators ramping limits, or time-integral constraints like those on run-of-river hydro generation or energy storage. For systems with relatively low levels of fluctuations, these limits are not frequently violated and it is not necessary incorporate them directly. However, higher levels of time-intermittent generation will create larger fluctuations and ramping events and the associated constraint violations will become more common. Standard primary and secondary controls are limited in their ability to balance these fluctuations, and better control design is needed to manage these larger fluctuations. Because these fluctuations are intimately connected to frequency deviations, they are of special concern because they may result in system-wide instabilities and loss of synchrony [3].

Other considerations for real-time power grid control systems are communication constraints and communication security[2]. Mechanisms that rely on central aggregation of the entire grid state followed by a centrally computed response will be vulnerable to communication failures and attacks on the communication network, making the overall system less robust. On the other hand, with significant renewable penetration, it is difficult to control a system purely based on local feedback, since under some conditions, it may be necessary to control distant generators in a correlated manner.

In this preliminary work, we explore a hybrid approach



Figure 1: Our model of the BPA Transmission Network

that combines the speed and security of fully distributed control with the extensive system visibility provided by centralized control. Our method performs a centralized lookahead dispatch that also computes optimal local feedback parameters for all controllable generation, thus enabling the system to respond to fluctuations based only on local observables. We expand our definition of local observables to include not just frequency but also real power flows to neighboring nodes. We use an ensemble of forecasts that capture various possible scenarios for the wind generation and loads over the next intra-dispatch period (5 min/15 min/ 1 hour) to design an optimal time-varying dispatch for all the generators, as well as local feedback functions that enable the generators to respond to fluctuations based on the local observables.

Our control design is split into 2 phases:

- a An off-line optimization phase where the distributed control gains are optimized jointly for the whole network in a central computer using extensive simulation of possible future wind generation and forecast scenarios. These gains are then communicated to each flexible resource (controllable device) in the transmission network. This off-line optimization would need to be re-run every time the statistics of possible future scenarios change significantly. In general, we expect this optimization to be run every time the generation re-dispatch changes.
- b An online response phase where each device imple-

ments its purely local control in response to local observables (local frequency, line flows etc.) on the pace of the standard primary controls.

We test our algorithm on historical data from the Bonneville Power Administration (BPA) system [4], an ideal test system for our algorithm as it has significant amounts of both hydro and wind generation. We show that our algorithm performs well, even in cases of significant wind ramps.

Our results (detailed in section 3) lead to the following important observations:

- a Local control based on response to frequency deviations and local line flows at each generation can keep frequency deviations down to the about 10 mHz while maintaining all the security and capacity constraints.
- b Proportional control on frequency deviations and feedback on line flows is sufficient. Adding a frequency-deviation integral response is unnecessary, which is advantageous because a distributed implementation of an integral term may cause instabilities due to errors in local frequency measurement, and also because it limits communication requirements.
- c Joint optimization of feedback parameters for frequency deviation and line flows is necessary. Independent optimization or removal of either term leads to poor control performance.
- d Optimization over a finite but representative set of future scenarios enables the generalization of the control to new unseen scenarios.

The rest of the paper is organized as follows: Section 2 describes the mathematical setting of the underlying control/optimization problem; we describe and discuss results of our numerical BPA experiments in Section 3; and Section 4 presents conclusions and explains our path forward.

2 Mathematical Formulation

2.1 Preliminaries

The power system is described by an undirected graph $\mathcal{G} = (\mathcal{E}, \mathcal{V})$ with edges \mathcal{E} and n vertices \mathcal{V} . The grid

is composed of loads (**l**), conventional generators (**g**) and renewable generators (**r**). The flexible resources in our grid are the online conventional generators **go**. We denote by \mathbf{p} the $|\mathcal{V}| \times 1$ vector of net active power injections at each node in the network and by $\mathbf{p}^{\mathbf{go}}, \mathbf{p}^{\mathbf{r}}, \mathbf{p}^{\mathbf{l}}$ the net active power injections due to online conventional generators, renewable generators and loads: $\mathbf{p} = \mathbf{p}^{\mathbf{go}} + \mathbf{p}^{\mathbf{r}} + \mathbf{p}^{\mathbf{l}}$. Each of these vectors is of size $|\mathcal{V}| \times 1$ with the convention that $\mathbf{p}_i^{\mathbf{go}} = 0$ or $\mathbf{p}_i^{\mathbf{r}} = 0$ if there is no conventional or renewable generator at node i . Note here that we make the assumption that there is only one generator or load at a given node. If there are multiple, we replace them by an equivalent single generator or load. For a vector \mathbf{v} with indices in \mathcal{V} , v_i denotes a particular component for $i \in \mathcal{V}$ and \mathbf{v}_S denotes the sub-vector $\{v_i : i \in S\}$ for a subset $S \subset \mathcal{V}$.

2.2 Scenarios

The control algorithm proceeds by analyzing an ensemble possible future scenarios and designs control strategies that optimize the system cost (defined below) across all scenarios. We define a scenario χ to be a collection of the following quantities:

- a Renewable generation over the time horizon of interest: $\mathbf{p}^{\mathbf{r}}(t)$.
- b Load profile over the time horizon of interest: $\mathbf{p}_0^{\mathbf{l}}(t)$.
- c A unit commitment (configuration of generators which are online, i.e. available for re-dispatch) **go**.

To define the control problem, we require a collection of scenarios Ξ and estimates for the probability of each scenario, i.e. $\Xi = \{\chi_i, \text{Prob}(\chi_i)\}$. We note that for a given collection Ξ , $\mathbf{p}^{\mathbf{r}}(t)$ and $\mathbf{p}_0^{\mathbf{l}}(t)$ (items a and b from above) will vary across the ensemble of scenarios, however, we take **go** (the unit commitment from c) fixed because we are designing the time-dependent dispatch and local feedback parameter for that particular **go**. In this work, we assume that the collection Ξ is finite. Typically, Ξ will be built up from load and wind forecasts from different forecasting methodologies weighted by confidences in each of these forecasts. Ξ could also include samples from a stochastic forecasting model based on climate models, historical data, meteorological sensors etc.

2.3 Control Formulation

We ignore electro-mechanical dynamical transients and work with a discrete-time quasi-static approximation of the system dynamics with fixed time step δ and integer time indices $t = 0, 1, \dots, T$: at each time step the power flows over lines are re-computed for configuration of consumption/generation at nodes evolving in discrete time. In general, the feedback can depend on any of the system variables, but we limit ourselves to local observables so that the control can be implemented in a completely distributed fashion at each generator after the dispatch and feedback parameters have been communicated.

For each generator $g \in \mathbf{g}^0$, we compute a time-varying dispatch $\mathbf{p}_g^0(t) : 0 \leq t \leq T$, proportional frequency response coefficient α_g^P , integral frequency response coefficient α_g^I , and a response coefficient to local flows $\{\alpha_{g \rightarrow i}^F : i \in \text{Neb}(g)\}$. Further, we denote by $\omega(t)$ the frequency deviation from the nominal frequency (50/60 Hz) at time t and by $\Omega(t)$ the integral of the frequency deviation, which in discrete-time is approximated by $\Omega(t) = \sum_{\tau=0}^t \gamma^{\tau-t} \omega(\tau)$ where $0 < \gamma < 1$ is a discount factor. In other words, the integral frequency term is simply a weighted sum of frequency deviations in the past, where frequency deviations that are further in the past receive a geometrically smaller weight. With the time varying dispatch and feedback parameters determined, the output of the generators is given by:

$$\mathbf{p}_g^{\mathbf{g}^0}(t) = \mathbf{p}_{0g}^{\mathbf{g}^0}(t) + \alpha_g^P \omega(t) + \alpha_g^I \Omega(t) + \sum_{i \in \text{Neb}(g)} \alpha_{g \rightarrow i}^F \mathbf{p}_{g \rightarrow i}(t).$$

Although our algorithm can incorporate nonlinear feedback, we choose feedback which is linear in the local observables for this initial work. In addition to generators, the real power consumption of loads responds to frequency changes, and we assume a simple linear load-frequency response given by

$$\mathbf{p}^l(t) = \mathbf{p}_0^l(t) + \beta^l \omega(t),$$

where the β^l are known from measurement where \mathbf{p}_0^l is the load at the nominal frequency (60 Hz). Combining the load and generator frequency response and the generators'

time varying dispatch, the system's equilibrium frequency is computed by enforcing power balance in the system:

$$\sum_{i \in \mathcal{V}} \mathbf{p}_i^{\mathbf{g}^0}(t) + \mathbf{p}_i^l(t) + \mathbf{p}_i^r(t) = 0 \implies \omega(t) = - \frac{\sum_i \mathbf{p}_{0i}^l(t) + \mathbf{p}_i^r(t) + \mathbf{p}_i^{\mathbf{g}^0}(t)}{\sum_i \beta_i^l}.$$

To compute power flow from the injections $\mathbf{p}(t) = \mathbf{p}^l(t) + \mathbf{p}^r(t) + \mathbf{p}^{\mathbf{g}^0}(t)$, we use a modified version of the DC Power flow equations based on a linearization of the AC Power flow equations around the nominal dispatch at the beginning of the control period $\mathbf{p}(0)$. The linearization gives us dynamic impedances $\mathbf{s}_{i \rightarrow j}^d$ that substitute for the line reactances in the DC power flow equations:

$$\mathbf{p}_i(t) = \mathbf{p}^0 + \sum_{j \in \text{Neb}(i)} \frac{\theta_i(t) - \theta_j(t)}{\mathbf{s}_{i \rightarrow j}^d},$$

$$\mathbf{p}_{i \rightarrow j}(t) = \mathbf{p}_{i \rightarrow j}^0 + \frac{\theta_i(t) - \theta_j(t)}{\mathbf{s}_{i \rightarrow j}^d}.$$

Such a linearization is reasonable assuming that the flow patterns do not change too much during the course of the control period.

In addition to several other constraints discussed below, we will also imposed a constraint on the total energy extracted from generators in the control period. Such constraints can represent the water discharge constraints on run-of-river hydro systems or state-of-charge constraints on energy storage devices. Therefore, we must also include the total energy extracted from each generator into the system state:

$$\mathbf{p}^l(t) = \sum_{\tau=0}^t \mathbf{p}^{\mathbf{g}^0}(\tau).$$

The overall system state consists of $\mathbf{x}(t) = [\Omega(t); \mathbf{p}^{\mathbf{g}^0}(t); \mathbf{p}^l(t)]$ (in Matlab notation), and the

system evolution can be summarized by:

$$\begin{aligned}
\omega(t) &= -\frac{\sum_i \mathbf{p}_{0i}^l(t) + \mathbf{p}_i^r(t) + \mathbf{p}_i^{\text{go}}(t)}{\sum_i \beta_i^l} \quad (1) \\
\Omega(t+1) &= \omega(t) + \gamma\Omega(t) \\
\mathbf{p}_g^{\text{go}}(t) &= \mathbf{p}_{0g}^{\text{go}}(t) + \alpha_g^P \omega(t) + \alpha_g^l \Omega(t) \\
&\quad + \sum_{i \in \text{Neb}(g)} \alpha_{g \rightarrow i}^F \mathbf{p}_{g \rightarrow i}(t) \\
\mathbf{p}^l(t) &= \mathbf{p}_0^l(t) + \beta^l \omega(t) \\
\mathbf{p}_i(t) &= \sum_{j \in \text{Neb}(i)} \frac{\theta_i(t) - \theta_j(t)}{\mathbf{s}_{i \rightarrow j}^d}, \mathbf{p}_{i \rightarrow j}(t) = \frac{\theta_i(t) - \theta_j(t)}{\mathbf{s}_{i \rightarrow j}^d}
\end{aligned}$$

2.4 Cost Functions

We consider a stochastic setting with many possible features, and it is unclear whether it is feasible to satisfy all constraints across all scenarios in Ξ . Therefore, we use a penalty function to enforce our constraints in a smooth manner. The penalty function has a magnitude of zero in a dead-band around the most feasible region and grows cubically with the magnitude of constraint violation:

$$\text{Pen}(a, l, u) = \begin{cases} 10^7((a-u)/(0.1 * (u+1)))^3 & \text{if } a \geq u \\ 10^7((l-a)/(0.1 * (l+1)))^3 & \text{if } a \leq l \\ 0 & \text{otherwise} \end{cases}.$$

Here, a is the value of the constrained quantity and l and u are the lower and upper bounds on a , respectively. We also adopt the convention that when a, l, u can be vectors (of the same size) and the penalty in this case is applied element-wise and added up. The penalty function is designed so that the resulting cost function is smooth (twice differentiable). However, if a violates the upper bound by 10%, a penalty of approximately 10^7 is incurred—a high enough penalty so that if a feasible solution exists across all scenarios, it will be found.

The cost function $\text{Cost}(\mathbf{x}(t), \mathbf{x}(t+1), t)$ is computed at each time step in the control period, but it requires state information from both t and $t+1$ so it can incorporate generator ramping limits. The cost includes seven terms that penalize both economic cost of supplying generation and deviations of the system state outside of normal operational bounds. The individual terms are:

1 Generation costs

$$\text{GenCost}(\mathbf{p}_{\text{go}}(t)) = \sum_{g \in \text{go}} c_{g1}(\mathbf{p}_g^{\text{go}})^2 + c_{g2}\mathbf{p}_g^{\text{go}} + c_{g3}.$$

2 Generation limit penalties

$$\text{Pen}(\mathbf{p}^{\text{go}}(t), \underline{\mathbf{p}}^{\text{go}}, \overline{\mathbf{p}}^{\text{go}}).$$

3 Ramping limit penalties

$$\text{Pen}\left(\underline{\mathbf{p}}_r^{\text{go}}, \frac{\mathbf{p}_r^{\text{go}}(t+1) - \mathbf{p}_r^{\text{go}}(t)}{\delta}, \overline{\mathbf{p}}_r^{\text{go}}\right).$$

4 Power flow thermal limit penalties

$$\sum_{i \rightarrow j \in \mathcal{E}} \text{Pen}(\mathbf{p}_{i \rightarrow j}(t), -\bar{\mathbf{p}}_{i \rightarrow j}, \bar{\mathbf{p}}).$$

5 Frequency deviation penalties

$$\text{Pen}(\omega(t), -0.01, 0.01)$$

6 Integral frequency deviation penalties

$$\text{Pen}(\Omega(t), -0.01, 0.01).$$

7 An integral deviation penalty on generation:

$$\text{Pen}(\mathbf{p}^l(T), 0.95\overline{E}^{\text{go}}, 1.05\overline{E}^{\text{go}})$$

Cost 1 simply represents the financial cost of energy from different generators. Costs 2-4 are normal power system constraints converted to costs using the penalty function defined above. Cost 5 is an additional penalty designed to constrain the system frequency to within a 10 mHz band, and Cost 6 is designed to constrain the deviation of the integral of the frequency deviation so that the frequency is not allowed to be low or high for extended periods of time. Finally, Cost 7 is designed to keep the total energy delivered by each controllable generator over the control period within a $\pm 5\%$ band around a $\mathbf{p}^l(T)$ mimicking constraint the would occur in either a run-of-river hydro system or an energy storage device.

2.5 Ensemble Optimal Control

The evolution equations listed in (1) are functions of a given scenario χ , therefore, we can think of the state as a function of the scenario χ and the control parameters $\alpha = \{\alpha^P, \alpha^I, \alpha^F, \mathbf{p}_0^{\text{go}}(t) : 0 \leq t \leq T\}$: $\mathbf{x}(\alpha, \chi, t)$. The overall optimization problem can then be written

$$\min_{\alpha} \sum_{\chi} \text{Prob}(\chi) \left(\sum_{\tau=0}^{T-1} \text{Cost}(\mathbf{x}(\alpha, \chi, \tau), \mathbf{x}(\alpha, \chi, \tau+1), \tau) \right)$$

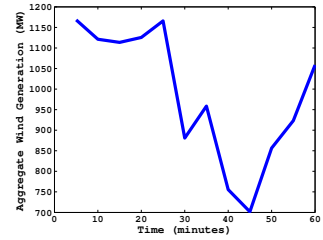
Subject to (1). (2)

We optimize this objective using a standard numerical optimization algorithm (LBFGS [5]). The gradients of the objective function can be computed efficiently using a forward propagation algorithm that uses the chain rule to propagate gradients in time. This computation can be easily vectorized over all the scenarios, leading to significant speedup if run on a cluster or on GPUs.

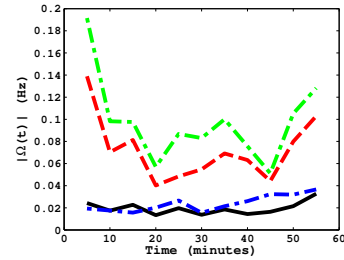
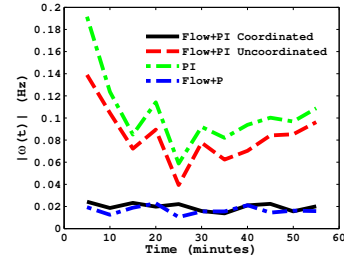
3 Numerical Results

3.1 Description of Test System

We test our algorithm using publicly available historical data for hydro and thermal generation, wind generation, and load from the Bonneville Power Administration (BPA) website[4]. We use a model of the BPA transmission system (shown in Fig. 1) that has 2209 buses and 2866 transmission lines. By identifying major hydroelectric stations on the transmission system and overlaying this onto a publicly available BPA wind site map [6], we located the existing wind farms on the BPA transmission system (as of January 2010). We located the meteorological stations where BPA collects wind data [7] in a similar manner. Using the same overlay, we used a simple incompressible air-flow model to infer hub height wind speeds at the wind farms. The resulting wind speeds were passed through the power curve of a standard 1.5-MW GE Wind Turbine which was scaled to the wind farm nameplate capacity to estimate the power output $\mathbf{p}^r(t)$ (in MW) at each wind farm as a function of time. When we aggregate our wind farm-specific estimates of wind generation, we typically over estimate the BPA aggregate data by 20%, which may caused by



(a) Aggregate Wind Generation



(b) Instantaneous and Integral Frequency Deviations

Figure 2: Comparison of control schemes. a) Aggregate wind generation from a period with significant ramping events. b) *Worst-case* frequency deviations over the control period for 18 validation scenarios not used in the control design.

several factors including: spilling of wind by BPA, under performance of wind farms relative to single-turbine estimates, or shortcomings in our model of interpolating wind speeds. BPA also provides aggregate load data [4] that we divide among the nodes in the network according to population densities. BPA also makes publicly available aggregate interchange flows [4], which we apportion to different tie lines in a similar manner.

To test our control algorithm on difficult conditions, we select a control period of one hour from 10:35 AM to 11:35 AM on February 12, 2010, when the wind generation was ramping significantly (shown in Fig. 2a). We then create 26 scenarios (site-specific wind profiles) for this period by adding random time-varying Gaussian noise to the wind speeds at each meteorological station (from which we infer site-specific wind generation as outlined above). We set the magnitude of the noise so as to match, on average, the aggregate wind generation hour-ahead forecast errors reported by BPA [8]. All the time series data used in our study was available at a 5-minute resolution.

Unit commitment data is missing from our model, therefore, we assume that all hydro generators larger than 300 MW are online and are all participating in frequency regulation. From inspection of the BPA historical generation data [4], we infer that the thermal generation dispatch is fixed over time. In our model, we replicate this dispatch by dividing the total thermal generation among the online thermal generators (randomly chosen).

3.2 Comparison of Various Control Schemes

For difficult wind ramping conditions, we illustrate the value of feedback based on local flows by comparing four control schemes. We use P to designate proportional control (to frequency deviations $\omega(t)$) and I designates integral control (to integral frequency deviations $\Omega(t)$). The control schemes we consider using are:

- 1 *PI*: Joint optimization of the time-varying dispatch $\mathbf{p}_0^{\text{go}}(t)$ and the local feedback parameters for $\omega(t)$ and $\Omega(t)$.

- 2 *Flow+PI Uncoordinated*: Time-varying dispatch $\mathbf{p}_0^{\text{go}}(t)$ plus feedback on $\omega(t)$, $\Omega(t)$ and local flows $\mathbf{p}_{g \rightarrow i}$ at each generator. The optimization in 1 is performed first followed by a second optimization over the flow feedback parameters.

- 3 *Flow+PI Coordinated*: Same as 2, but the optimization is performed jointly.

- 4 *Flow+P*: Same as 3, but without feedback on $\Omega(t)$.

The experimental protocol is as follows. We setup each of the four optimization problems according to Eqs. 2 with the scenarios described in Section 3.1 and determine a single set of feedback parameters for each of the four feedback schemes. We use 8 of the 26 created scenarios as input to the optimization algorithm. The remaining 18 unseen scenarios are reserved for validation of the control policy discovered by the optimization algorithm. We note that all four control strategies are able to achieve similar generation costs while maintaining all the other constraints (line thermal capacities, ramping limits, and integral energy constraints), however, there are significant differences in the quality of the frequency regulation. Figure 2b shows the worst-case frequency deviations over the 18 validation scenarios. The frequency deviations are at an unacceptable level (.1-.2 Hz) when using just PI feedback (scheme 1). If the flow feedback is included but optimized separately (scheme 2), there is little improvement. However, if the PI and flow feedback are coordinated via joint optimization (scheme 3), the frequency deviations are reduced to an acceptable level. Interestingly, removing the feedback on the integral of the frequency deviations (scheme 4) does not impact the frequency deviations significantly relative to scheme 3.

3.3 Discussion of the Results

The distributed frequency control method we have presented benefits greatly from the incorporation of local power flows as demonstrated in Fig. 2b. There are several possible reasons for this improved performance. First, power flows make the local generation-load imbalances visible to the generators so that the closest generators respond, effectively screening the more distant generators from the need to respond. When compared to feedback based on frequency deviation, which is a global measure

of the imbalance, feedback on local power flows confines imbalances to shorter spatial scales with a corresponding decrease in the time scale of the response. An alternative explanation is that the optimization over the ensemble of possible futures in Eq. 2 is acting as a sort of machine learning that encodes correlations between the wind prediction errors and the resulting local power flows into the flow feedback parameters. When wind prediction error occurs, the change in power flows drives the feedback to nearly compensate for the error without a frequency deviation existing for any significant length of time. More numerical experiments are required to distinguish between these two (and other) possibilities. In both of the possibilities discussed above, variations in the local power flows appear to be acting as “pseudo-communication” channels between the renewable and controllable generators. Such a communication analogy may help explain why the independent optimizations in scheme 2 does not yield significant improvement in control performance. The first optimization over frequency deviations may effectively washout the important local information in the power flows such that it is not available when optimizing over power flows.

4 Conclusions and Future Work

We introduced a control architecture based on off-line centralized optimization that can occur on a slow time scale coupled that sets the feedback parameters for fast distributed control of generation. The control scheme takes into account explicitly the variability in renewable generation using ensemble control. We showed that local feedback based on line flows and frequency deviations is sufficient to maintain all operational constraints and limit frequency deviations to an acceptable level even when the system is experiencing significant ramps in wind generation. Our method exploits the hour-scale predictability of wind energy while using the off-line optimization to re-adjust control policies over longer timescales where wind predictability suffers. Our hybrid approach has the potential enable even higher levels time-intermittent renewable generation than presented here, and it can do so without real-time computation or communication.

These results are quite exciting and promising, however, they are preliminary and much work needs to be done to ensure the viability of this scheme in practice.

- Dynamical simulations are needed to check the dynamical stability of a grid with flow feedback. If these simulations show that the scheme is unstable, we believe that this can be rectified by appropriate exciter control at the generators to damp the fast electro-mechanical transients.
- The scenario approach can be extended to include the (N-1) security criterion, so that the optimized control strategy can deal with contingencies arising from the failure of a grid component.
- It is possible that flow feedback acts as a pseudo-communication channel between generators in the absence of a dedicated communication channel. It would be interesting to investigate this from an information theoretic point of view and investigate how much of information can be encoded in the flows.
- We have used the simplest possible algorithmic approach by defining a smooth version of the optimization problem using penalty functions solving it using a generic LBFGS algorithm [5]. Second-order algorithms such as Stagewise Newton[9] or Differential Dynamic Programming (DDP)[10] efficiently exploit the problem structure of deterministic optimal control problems. These can be leveraged in our ensemble control context by noting that when the feedback parameters $\alpha^I, \alpha^P, \alpha^F$ are fixed, we have a deterministic optimal control problem in $\mathbf{p}_0^{\text{go}}(t)$ for each scenario. We have also been working on a Gauss-Newton algorithm for optimizing the fixed feedback $\alpha^I, \alpha^P, \alpha^F$ efficiently. One can perform alternate minimization of $\mathbf{p}_0^{\text{go}}(t)$ and $\alpha^I, \alpha^P, \alpha^F$ to get an efficient algorithm for optimizing both. Further, we note that when feedback does not include the integral term $\Omega(t)$, the ensemble control problem is a convex programming problem, and the global optimum can be found efficiently using specialized convex optimization techniques.
- We plan to incorporate more accurate AC modeling of power flows taking advantage of most recent

advances in analysis and algorithms related to optimizations of nonlinear power flows, e.g. [11, 12].

- The integral energy constraint we introduced can also model energy storage, and our algorithm can easily be extended to incorporate distributed control of energy storage.

Acknowledgments

The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396. SB and MC acknowledge partial support of NMC via NSF collaborative research grant ECCS-1128325 on “Power Grid Spectroscopy”.

References

- [1] P. Kundur. *Power system stability and control*. McGraw-Hill, 1994.
- [2] K. Tomsovic, D.E. Bakken, V. Venkatasubramanian, and A. Bose. Designing the next generation of real-time control, communication, and computations for large power systems. *Proceedings of the IEEE*, 93(5):965–979, 2005.
- [3] J. Eto, J. Undrill, P. Mackin, H. Illian, A. Martinez, M. O’Malley, and K. Coughlin. Use of frequency response metrics to assess the planning and operating requirements for reliable integration of variable renewable generation. LBNL-4142E; <http://transmission.bpa.gov/business/operations/wind/>, 2010.
- [4] BPA wind data. <http://transmission.bpa.gov/business/operations/wind/>.
- [5] M. Schmidt. minfunc., 2005.
- [6] BPA wind site map. http://transmission.bpa.gov/PlanProj/Wind/documents/BPA_wind_map_2011.pdf.
- [7] BPA meteorological site map. <http://transmission.bpa.gov/business/operations/wind/MetData.aspx>.
- [8] BPA wind forecast data. <http://transmission.bpa.gov/Business/Operations/Wind/forecast/forecast.aspx>.
- [9] J.F.A.D.O. Pantoja. Algorithms for constrained optimization problems. *Differential dynamic programming and Newton’s method. International Journal of Control*, 47:1539–1553, 1983.
- [10] D.H. Jacobson. Second-order and second-variation methods for determining optimal control: A comparative study using differential dynamic programming. *International Journal of Control*, 7(2):175–196, 1968.
- [11] J. Lavaei and S. Low. Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems*, 27(1):92–107, 2012.
- [12] M. Kraning, E. Chu, J. Lavaei, and S. Boyd. Message passing for dynamic network energy management. http://www.stanford.edu/~boyd/papers/decen_dyn_opt.html, 2012.